

CLAIMS

1. A method of transferring a packet to a computer system, wherein the packet is received at a communication device from a network, comprising:

parsing a header portion of a first packet received at a communication device to determine if said first packet conforms to a pre-selected protocol;

generating a flow key to identify a first communication flow that includes said first packet;

transferring said first packet to a host computer system for processing in accordance with said pre-selected protocol; and

associating an operation code with said first packet, wherein said operation code indicates a status of said first packet.

2. The method of claim 1, wherein said parsing comprises:

copying a header portion of said first packet into a header memory; and

examining said header portion according to a series of parsing instructions;

wherein said parsing instructions are configured to reflect a set of pre-selected communication protocols.

3. The method of claim 2, wherein said parsing instructions are updateable.

4. The method of claim 2, further comprising copying a value from a field in a header of said header portion.

5. The method of claim 1, wherein said parsing comprises:

extracting an identifier of a source of said first packet from said header portion;

and

extracting an identifier of a destination of said first packet from said header portion.

6. The method of claim 5, wherein said generating comprises combining said source identifier and said destination identifier.
7. The method of claim 1, wherein said generating comprises retrieving an identifier of a communication connection from said header portion.
8. The method of claim 1, further comprising storing said first packet in a packet memory prior to said transferring.
9. The method of claim 1, further comprising storing said flow key in a flow database, wherein said flow database is configured to facilitate management of said first communication flow.
10. The method of claim 9, further comprising associating a flow number with said first packet, wherein said flow number comprises an index of said flow key within said flow database.
11. The method of claim 10, further comprising storing said flow number in a flow memory.
12. The method of claim 9, further comprising updating an entry in said flow database associated with said flow key when a second packet in said first communication flow is received.
13. A computer readable storage medium storing instructions that, when executed by a computer, cause the computer to perform a method of transferring a packet received at a network interface from a network to a host computer system, the method comprising:
 - receiving a packet from a network at a network interface for a host computer system;
 - parsing a header portion of said packet to extract an identifier of a source entity and an identifier of a destination entity;

generating a flow key from said source identifier and said destination identifier to identify a communication flow comprising said packet;

determining whether a header in said header portion conforms to a pre-selected protocol;

storing said flow key in a database;

associating an operation code with said packet, wherein said operation code identifies a status of said packet;

storing said packet in a packet memory;

if said header conforms to said pre-selected protocol:

storing a data portion of said packet in a re-assembly buffer; and

storing said header portion in a header buffer; and

if said header conforms to a protocol other than said pre-selected protocol, storing said packet in a non-re-assembly buffer.

14. The method of claim 1, wherein said associating comprises:

retrieving one or more header fields of said header portion; and analyzing said header fields to determine said status of said first packet.

15. The method of claim 14, wherein said analyzing comprises:

determining whether said first packet includes a data portion; and

if said first packet includes a data portion, determining whether said data portion exceeds a pre-determined size.

16. The method of claim 14, wherein said analyzing comprises determining whether said first packet was received out of order in said first communication flow.

17. The method of claim 1, further comprising storing said operation code in a control memory.

18. The method of claim 1, wherein said first packet is determined to conform to said pre-selected protocol, said transferring comprising:

storing a data portion of said first packet in a re-assembly storage area, wherein said re-assembly storage area is configured to only store data portions of packets in said first communication flow; and

storing one or more headers from said header portion in a header storage area.

19. The method of claim 1, wherein said transferring comprises:

if said first packet is smaller than a predetermined threshold, storing said first packet in a first storage area; and

if said first packet is larger than said predetermined threshold, storing said first packet in a second storage area.

20. The method of claim 1, further comprising determining whether a second packet received from said network is part of said first communication flow.

21. The method of claim 20, wherein said determining comprises:

maintaining a packet memory configured to store one or more packets received from said network;

maintaining a flow memory configured to store, for each of said one or more packets, an identifier of a communication flow comprising said packet; and

searching said flow memory for a first identifier of said first communication flow.

22. The method of claim 21, wherein said first identifier comprises said flow key.

23. The method of claim 21, wherein said first identifier comprises a flow number of said first packet, wherein said flow number is an index of said flow key within a flow database.

24. A method of transferring a packet to a computer system, wherein the packet is received at a communication device from a network, comprising:

parsing a header portion of a first packet received at a communication device to determine if said first packet conforms to a pre-selected protocol;

 generating a transmit control block (TCB) to identify a first communication flow that includes said first packet;

 transferring said first packet to a host computer system for processing in accordance with said pre-selected protocol; and

 associating a summary with said first packet, wherein said summary indicates a status of said first packet.

25. The method of claim 24, further comprising:

transferring said TCB from said host computer system to said communication device.

26. The method of claim 24, further comprising:

 receiving a second packet from a second communication flow; and

 processing said second packet by said communication device in accordance with said TCB.

27. The method of claim 1, further comprising alerting said host computer system to the arrival of said first packet.

28. The method of claim 1, further comprising:

 maintaining a packet memory configured to store packets received from said network; and

 randomly discarding a packet if said packet memory contains a pre-determined level of traffic.

29. The method of claim 1, wherein said parsing includes determining, by hardware of the communication device, a session layer protocol of the header.

30. The method of claim 1, wherein said generating a flow key includes initializing a communication control block (CCB) during Transport Control Protocol (TCP) connection setup.

31. The method of claim 1, wherein said communication device is a network interface.

32. A method of transferring a packet received at a network interface to a host computer system, comprising:

- receiving a packet from a network;
- storing said packet in a packet memory;
- parsing a header portion of said packet;
- extracting a value stored in said header portion;
- identifying a communication flow comprising said packet;
- determining whether a header in said header portion conforms to a pre-selected protocol;

- determining whether a second packet in said packet memory is part of said communication flow;

- if the host computer system contains a plurality of processors, identifying a processor to process said packet; and

- storing said packet in a host memory area.

33. A method of transferring a packet received at a network interface from a network to a host computer system, comprising:

- receiving a packet from a network at a network interface for a host computer system;

- parsing a header portion of said packet to extract an identifier of a source entity and an identifier of a destination entity;

- generating a flow key from said source identifier and said destination identifier to identify a communication flow comprising said packet;

- determining whether a header in said header portion conforms to a pre-selected protocol;

storing said flow key in a database;
associating an operation code with said packet, wherein said operation code identifies a status of said packet;
storing said packet in a packet memory;
if said header conforms to said pre-selected protocol:
storing a data portion of said packet in a re-assembly buffer; and
storing said header portion in a header buffer; and
if said header conforms to a protocol other than said pre-selected protocol, storing said packet in a non-re-assembly buffer.

34. The method of claim 33, wherein said parsing comprises executing a series of updateable instructions configured to parse a packet header conforming to one of a set of pre-selected protocols.

35. The method of claim 33, further comprising storing said operation code in a control memory.

36. The method of claim 33, further comprising storing a flow number of said packet in a flow memory, wherein said flow number comprises an index of said flow key in said database.

37. The method of claim 36, further comprising indicating whether said packet memory includes another packet with said flow number or said flow key.

38. The method of claim 33, wherein the host computer system comprises multiple processors, further comprising identifying a first processor in the host computer system to process said packet in accordance with said pre-selected protocol.

39. The method of claim 38, further comprising:
receiving a second packet at said network interface, wherein said second packet is part of a second communication flow; and

identifying a second processor in the host computer system to process said second packet.

40. The method of claim 33, further comprising informing said host computer system of said receipt of said packet.

41. The method of claim 1, wherein said transferring said first packet to a host computer system comprises:

analyzing a session layer part of said header to allocate a list of memory addresses of the host computer for storing data from a plurality of packets that correspond to said communication flow.

42. An apparatus for transferring a packet to a host computer system, comprising:

a traffic classifier configured to classify a first packet received from a network by a communication flow that includes said first packet;

a packet memory configured to store said first packet;

a packet batching module configured to determine whether another packet in said packet memory belongs to said communication flow; and

a flow re-assembler configured to re-assemble a data portion of said first packet with a data portion of a second packet in said communication flow;

wherein said first packet data portion and said second packet data portion are stored in a host computer memory area to enable efficient transfer of said memory area contents.

43. The apparatus of claim 42, wherein said traffic classifier comprises:

a parser configured to parse a header portion of said first packet;

a flow database configured to store a flow key identifying said communication flow; and

a flow database manager configured to manage said flow database;

wherein said flow key is generated from an identifier of a source of said first packet and an identifier of a destination of said first packet.

44. A computer system for receiving a packet from a network, comprising:

- a memory configured to store packets received from a network; and
- a communication device configured to receive a first packet from said network,

the communication device comprising:

- a parser configured to extract information from a header portion of a first packet;
- a flow manager configured to examine said information;
- a flow database configured to store an identifier of a first communication flow comprising multiple packets, including said first packet; and
- a re-assembler for storing data portions of said multiple packets in a first portion of said memory; and
- a processor for processing said first packet.

45. The apparatus of claim 42, further comprising:

- a load distributor for identifying a first processor within the host computer system for processing said first packet and said second packet;

wherein said load distributor identifies a second processor in the host computer system for processing a packet from a different communication flow.

46. The method of claim 1, wherein said operation code indicates whether the packet corresponds to Transport Control Protocol (TCP).

47. A device for receiving a packet from a network and transferring the packet to a host computer system, comprising:

- a parser configured to parse a header portion of a packet received from a network, wherein said parsing comprises:

determining whether a header within said header portion conforms to one of a set of communication protocols; and

if said header conforms to one of said communication protocols, extracting information from said header portion to identify a communication flow to which said packet belongs;

a flow memory configured to store a flow identifier for identifying said communication flow;

a flow manager configured to assign an operation code to said packet, wherein said operation code:

indicates a status of said packet; and

indicates a manner of transferring said packet to the host computer system;

a packet memory configured to store said packet; and

a transfer module configured to transfer said packet from said packet memory to a host computer system in accordance with said operation code.

48. The device of claim 47, wherein the device is a network interface.

49. The device of claim 47, said flow memory comprising a flow database configured to store a flow key, wherein said flow key is assembled from an identifier of a source of said packet and an identifier of a destination of said packet.

50. The device of claim 47, wherein said flow manager is further configured to update said flow memory as additional packets in said communication flow are received from the network.

51. The device of claim 47, said flow memory comprising a flow memory configured to store a flow number, wherein said flow number comprises an index of said communication flow in a flow database.

52. The device of claim 47, further comprising a control memory configured to store said operation code.

53. A computer readable storage medium storing instructions that, when executed by a computer, cause the computer to perform a method of transferring a packet from a communication device to a host computer, the method comprising: parsing a header portion of a first packet received at a communication device to determine if said first packet conforms to a pre-selected protocol; generating a flow key to identify a first communication flow that includes said first packet; transferring said first packet to a host computer system for processing in accordance with said pre-selected protocol; and associating an operation code with said first packet, wherein said operation code indicates a status of said first packet.

54. The device of claim 47, wherein said host computer system is a multi-processor host computer system, further comprising a load distributor configured to select one of said multiple processors for processing said packet in accordance with one of said communication protocols.

55. The device of claim 47, wherein said transfer module is configured to transfer a data portion of said packet into one of a set of host memory areas in accordance with said operation code.

56. The device of claim 47, further comprising a packet batching module configured to determine whether said packet memory contains another packet in said communication flow.

57. The method of claim 1, further comprising:
receiving, by said communication device, a second packet that corresponds to said communication flow, said second packet including at least a transport layer header and data; and
transferring said data from said communication device to said host computer system, without transferring said transport layer header from said communication device to said host computer system.

58. An apparatus for transferring a packet from a network to a host computer system, comprising:

 a parser module configured to:

 parse a header portion of a first packet received from a network to extract an identifier of a source of said first packet and an identifier of a destination of said first packet;

 generate a flow key from said source identifier and said destination identifier to identify a communication flow comprising said first packet; and

 determine whether a header in said header portion conforms to a pre-selected protocol;

 a flow database configured to store said flow key;

 a flow database manager configured to associate an operation code with said first packet, wherein said operation code identifies a status of said first packet;

 a packet memory configured to store said first packet; and

 a transfer module configured to:

 if said header conforms to said pre-selected protocol:

 store a data portion of said first packet in a re-assembly buffer; and

 store said header portion in a header buffer; and

 if said header conforms to a protocol other than said pre-selected protocol,

 store said packet in a non-re-assembly buffer.

59. The apparatus of claim 58, wherein said transfer module comprises a re-assembly engine configured to re-assemble, in said re-assembly buffer, a data portion of said first packet with a data portion of a second packet in said first communication flow.

60. The apparatus of claim 58, further comprising a flow memory configured to store a flow number associated with said first packet, wherein said flow number comprises an index of said flow key in said flow database.

61. The apparatus of claim 58, further comprising:

a load distributor configured to identify a first processor in said host computer system for processing said first packet, said first processor being identified on the basis of said flow key;

wherein said host computer system is a multi-processor computer system; and
wherein a second processor in said host computer system is identified for processing a packet from a communication flow other than said first communication flow.

62. The apparatus of claim 58, further comprising:

a packet batching module configured to determine whether said packet memory includes another packet in said first communication flow.